arXiv:1412.0962v2 [cs.CG] 3 Dec 2014

Batched Point Location in SINR Diagrams via Algebraic Tools*

Boris Aronov† Matthew J. Katz‡

December 4, 2014

Abstract

The *SINR model* for the quality of wireless connections has been the subject of extensive recent study. It attempts to predict whether a particular transmitter is heard at a specific location, in a setting consisting of n simultaneous transmitters and background noise. The SINR model gives rise to a natural geometric object, the *SINR diagram*, which partitions the space into n regions where each of the transmitters can be heard and the remaining space where no transmitter can be heard.

Efficient *point location* in the SINR diagram, i.e., being able to build a data structure that facilitates determining, for a query point, whether any transmitter is heard there, and if so, which one, has been recently investigated in several papers. These planar data structures are constructed in time at least quadratic in n and support logarithmic-time approximate queries. Moreover, the performance of some of the proposed structures depends strongly not only on the number n of transmitters and on the approximation parameter ε, but also on some geometric parameters that cannot be bounded *a priori* as a function of n or ε.

In this paper, we address the question of *batched* point location queries, i.e., answering many queries simultaneously. Specifically, in one dimension, we can answer n queries *exactly* in amortized polylogarithmic time per query, while in the plane we can do it approximately.

In another result, we show how to answer n^2 queries *exactly* in amortized polylogarithmic time per query, assuming the queries are located on a possibly non-uniform $n \times n$ grid.

All these results can handle *arbitrary* power assignments to the transmitters. Moreover, the amortized query time in these results depends only on n and ε.

Finally, these results demonstrate the (so far underutilized) power of combining algebraic tools with those of computational geometry and other fields.

*Work on this paper by B.A. has been partially supported by NSF Grants CCF-11-17336 and CCF-12-18791. Work on this paper by M.K. has been partially supported by grant 1045/10 from the Israel Science Foundation.

†Department of Computer Science and Engineering, Polytechnic School of Engineering, New York University, Brooklyn, NY 11201, USA; boris.aronov@nyu.edu.

‡Department of Computer Science, Ben-Gurion University, Israel; matya@cs.bgu.ac.il.

1 Introduction

The *SINR (Signal to Interference plus Noise Ratio) model* attempts to more realistically predict whether a wireless transmission is received successfully, in a setting consisting of multiple simultaneous transmitters in the presence of background noise. In particular, it takes into account the attenuation of electromagnetic signals. The SINR model has been explored extensively in the literature [18].

Let $\mathcal{S} = \{s_1, \ldots, s_n\}$ be a set of n points in the plane representing n transmitters. Let $p_i > 0$ be the transmission power of transmitter s_i, $i = 1, \ldots, n$. In the *SINR model*, a receiver located at point q is able to receive the signal transmitted by s_i if the following inequality holds:

$$\frac{\frac{p_i}{|q-s_i|^\alpha}}{\Sigma_{j\neq i}\frac{p_j}{|q-s_j|^\alpha} + N} \geq \beta\,,$$

where $|a - b|$ denotes the Euclidean distance between points a and b, and $\alpha > 0$, $\beta > 1$, and $N > 0$ are given constants (N represents the background noise). This inequality is also called the *SINR inequality*, and when it holds, we say that q *receives* (or *hears*) s_i; we refer to the left hand side of the inequality as *SIN ratio* (for receiver q w.r.t. transmitter s_i).

Notice that, since $\beta > 1$, a necessary condition for q to receive s_i is that $p_i/|q - s_i| > p_j/|q - s_j|$, for any $j \neq i$. In particular, in the *uniform power setting* where $p_1 = p_2 = \cdots = p_n$, a necessary condition for q to receive s_i is that s_i is the closest to q among the transmitters in $\mathcal{S}$. This simple observation implies that, for any point q in the plane, either exactly one of the transmitters is received by q or none of them is. Thus, one can partition the plane into n not necessarily connected reception regions R_i, one per transmitter in $\mathcal{S}$, plus an additional region $R_\emptyset$ consisting of all points where none of the transmitters is received. This partition is called the *SINR diagram* of $\mathcal{S}$. Consider the *multiplicatively-weighted Voronoi diagram* D of $\mathcal{S}$ in which the region V_i associated with s_i consists of all points q in the plane for which $\frac{1}{\sqrt[\alpha]{p_i}}|q-s_i| < \frac{1}{\sqrt[\alpha]{p_j}}|q-s_j|$, for any $j \neq i$ (see Section B.2). Then $R_i \subset V_i$.

In a seminal paper, Avin et al. [6] studied properties of SINR diagrams, focusing on the uniform power setting. Their main result is that in this setting the reception regions R_i are convex and fat. (Here, R_i is *fat* if the ratio between the radii of the smallest disk centered at s_i containing R_i and the largest disk centered at s_i contained in R_i is bounded by some constant.) In the non-uniform power setting, on the other hand, the reception regions are not necessarily connected, and their connected components are not necessarily convex or fat. In fact, they may contain holes [16].

A natural question that one may ask is: "Given a point q in the plane, does q receive one of the transmitters in $\mathcal{S}$, and if yes which one?" Or equivalently: "Which region of the SINR diagram does q belong to?" The latter question is referred to as a *point-location query* in the SINR diagram of $\mathcal{S}$. We can answer it in linear time by first finding the sole candidate, s_i, as the transmitter for which the ratio $\frac{1}{\sqrt[\alpha]{p}}|q - s|$ is minimum, and then evaluating the SIN ratio and comparing it to β. To facilitate multiple queries, one may want to build a data structure that can guarantee faster response. We can expedite the first step by constructing the appropriate Voronoi diagram $D = D(\mathcal{S})$ together with a point-location structure, so that the sole candidate transmitter for a point q can be found in $O(\log n)$ time. However, the boundary of the region R_i is described by a degree-$\Theta(n)$ algebraic curve; it seems difficult (impossible, in general?) to build a data structure that can quickly determine the side of the curve a given point lies on. The answer is not even obvious in one dimension (where the transmitters and potential receivers all lie on a line), as there R_i is a collection of intervals delimited by roots of a polynomial of degree $\Theta(n)$.

The problem has been approached by constructing data structures for *approximate* point location in SINR diagrams. All approaches use essentially the same logic: first find the sole candidate s_i that the query point q may hear and then approximately locate q in R_i. This is done by constructing two sets R_i^+, R_i^- such that $R_i^+ \subset R_i \subset R_i^- \subset V_i$,[1] and preprocessing them for point location. In the region R_i^+ reception of s_i is guaranteed, so if $q \in R_i^+$, return "can hear s_i." Outside of R_i^- one cannot hear s_i, so if $q \notin R_i^-$, return "cannot hear anything." The set $R_i^- \setminus R_i^+$ is where the approximation occurs: s_i may or may not be heard there, so if $q \in R_i^-$ but $q \notin R_i^+$, return "may or may not hear s_i."

Two different notions of approximation have appeared in the literature. In the first [6, 16], it is guaranteed that the uncertain answer is only given infrequently, namely that $area(R_i^- \setminus R_i^+) \leq \varepsilon \cdot area(R_i)$, for a suitable parameter $\varepsilon > 0$. In the second [16], it is promised that the SIN ratio for every point in $R_i^- \setminus R_i^+$ lies within $[c_1\beta, c_2\beta]$ for suitable constants c_1, c_2 with $0 < c_1 < 1$, $c_2 > 1$.

We now briefly summarize previous work. Observing the difficulty of answering point-location queries exactly, Avin et al. [6] resorted to approximate query answers in the *uniform power* setting. Given an $\varepsilon > 0$ they build a data structure in total time $O(n^2/\varepsilon)$ and space $O(n/\varepsilon)$ that can be wrong only in a region of area $\varepsilon \cdot area(R_i)$ for each s_i (i.e., approximation of the first type described above). It supports logarithmic-time queries.

In a subsequent paper, Kantor et al. [16] studied properties of SINR diagrams in the *non-uniform power* setting. After revealing several interesting and useful properties, such as that the reception regions in the $(d+1)$-dimensional SINR diagram of a d-dimensional scene are connected, they present several solutions to the problem of efficiently answering point-location queries. One of them uses the second type of approximation, with $c_1 = (1-\varepsilon)^{2\alpha}$ and $c_2 = (1+\varepsilon)^{2\alpha}$, for a prespecified $\varepsilon > 0$. Queries can be performed in time $O(\log(n \cdot \varphi/\varepsilon))$, where φ is an upper bound on the fatness parameters of the reception regions (which cannot be bounded as a function of n or ε). The size of this data structure is $O(n \cdot \varphi'/\varepsilon^2)$ and its construction time is $O(n^2 \cdot \varphi'/\varepsilon^2)$, where $\varphi' > \varphi^2$ is some function of the fatness parameters of the reception regions.

Although highly non-trivial, the known results for point location in the SINR model are unsatisfactory, in that they suffer from very large preprocessing times. Moreover, in the non-uniform setting, the bounds include geometric parameters such as φ and φ' above, which cannot be bounded as a function of n or ε. In this paper we focus on *batched* point location in the SINR model. That is, given a set $\mathcal{Q}$ of m query points, determine for each point $q \in \mathcal{Q}$ whether it receives one of the transmitters in $\mathcal{S}$, and if yes, which one. Often the set of query points is known in advance, for example, in the planning stage of a wireless network or when examining an existing network. In these cases, one would like to exploit the additional information to speed up query processing. We achieve this goal in the SINR model; that is, we devise efficient approximation and exact algorithms for batched point location in various settings. Our algorithms use a novel combination of sophisticated geometric data structures and tools from computer algebra for multipoint evaluation, interpolation, and fast multiplication of polynomials and rational functions. For example, consider 1-dimensional batched point location where $m = n$ and power is non-uniform. We can answer *exactly* a point-location query in amortized time $O(\log^2 n \log\log n)$. Considering the same problem in the plane, for any $\varepsilon > 0$, we can approximately answer a query in amortized time polylogarithmic in n and ε, as opposed to the result of Kantor et al. [16] mentioned above in which the bounds depend on additional geometric parameters which cannot be bounded as a function of n or ε.

[1] Notice that we have not followed the original notation in the literature, for consistency with our notation below.

1.1 Related work

The papers most relevant to ours are those by Avin et al. [6] and Kantor et al. [16] discussed above. Avin et al. [5] also considered the problem of handling queries of the following form (in the uniform-power setting): Given a transmitter s_i and query point q, does q receive s_i by successively applying interference cancellation? (Interference cancellation is a technology that enables a point q to receive a transmitter s, even if s's signal is not the strongest one received at q; see [5] for further details.)

Gupta and Kumar [11] initiated an extensive study of the *maximum capacity* and *scheduling* problems in the SINR model. Given a set L of sender-receiver pairs (i.e., directional links), the *maximum capacity* problem is to find a *feasible* subset of L of maximum cardinality, where $L' \subseteq L$ is *feasible* if, when only the senders of the links in L' are active, each of the links in L' is feasible according to the SINR inequality. The *scheduling* problem is to partition L into a minimum number of feasible subsets (i.e., rounds). We mention several papers and results dealing with the maximum capacity and scheduling problems. Goussevskaia et al. [10] showed that both problems are NP-complete, even in the uniform power setting. Goussevskaia et al. [9], Halldórsson and Wattenhofer [14], and Wan et al. [26] gave constant-factor approximation algorithms for the maximum-capacity problem yielding an $O(\log n)$-approximation algorithm for the scheduling problem, assuming uniform power. In [9] they note that their $O(1)$-approximation algorithm also applies to the case where the ratio between the maximum and minimum power is bounded by a constant and for the case where the number of different power levels is constant. More recently, Halldórsson and Mitra [13] have considered the case of oblivious power. This is a special case of non-uniform power where the power of a link is a simple function of the link's length. They gave an $O(1)$-approximation algorithm for the maximum capacity problem, yielding an $O(\log n)$-approximation algorithm for scheduling. Finally, the version where one assigns powers to the senders (i.e., with power control) has also been studied, see, e.g., [2, 12, 13, 17, 21].

1.2 Our tools and goals

Besides making progress on the actual problems being considered here, we view this work as another demonstration of what we hope to be a developing trend of combining tools from the computer algebra world with those of computational geometry and other fields. Several relatively recent representatives of such synergy show examples of seemingly impossible speed-ups in geometric algorithms by expressing a subproblem in algebraic terms [1, 19, 20]. The algebraic tools themselves are mostly classical ones, such as Fast Fourier Transform, fast polynomial multiplication, multipoint evaluation, and interpolation; see Appendix A for details. We combine them with only slightly newer tools from computational geometry, such as Voronoi diagrams, point location structures in the plane, fast exact and approximate nearest-neighbor query data structures, and range searching data structures; refer to Appendix B. One very recent result we need is that of Har-Peled and Kumar [15] that, as a special case, allows one to build a compact data structure for approximating multiplicatively weighted nearest-neighbor queries in the plane; the exact version appears to require building the classical multiplicatively weighted Voronoi diagram, which is a quadratic-size object.

We hope that the current work will lead to further productive collaborations between computational geometry and computer algebra.

1.3 Our results

We now summarize our main results. We use O^* notation to suppress logarithmic factors and O_ε to denote polynomial dependence on $1/\varepsilon$, where $\varepsilon > 0$ is the approximation parameter. In general, we present algorithms for both the uniform-power and non-uniform-power settings, where the algorithms of the former type are usually somewhat simpler.

- In one dimension, we can perform n queries among n transmitters exactly in $O^*(n)$ total time; see Section 2.
- In two dimensions, we can perform n queries among n transmitters approximately in $O^*_\varepsilon(n)$ total time; see Section 3.4.
- We can also facilitate exact batched queries when they form a grid; see Sections 3.3 for the exact statement.

2 Batched point location on the line

In this section $\mathcal{S}$ is a set of $n \geq 3$ point transmitters and $\mathcal{Q}$ is a set of m query points, both on the line. We first consider the *uniform-power version* of the problem, where each transmitter has transmission power 1 (i.e., $p_1 = \cdots = p_n = 1$), and then extend the approach to the arbitrary power version.

2.1 Uniform power

A query point q receives s_i if and only if

$$\frac{\frac{1}{|q-s_i|^\alpha}}{\Sigma_{j\neq i}\frac{1}{|q-s_j|^\alpha} + N} \geq \beta\,.$$

Recall that, since $\beta > 1$, if q receives one of the transmitters, then it must be the transmitter that is closest to it; we call it the *candidate* transmitter for q and denote it by $s(q) = s(q, \mathcal{S})$.

Next, we define a univariate function f as

$$f(q) \coloneqq \sum_{j=1}^{n} \frac{1}{|q - s_j|^\alpha}\,.$$

Then, q can hear its candidate transmitter $s(q)$ if and only if

$$E(q) \coloneqq \frac{\frac{1}{|q-s(q)|^\alpha}}{f(q) - \frac{1}{|q-s(q)|^\alpha} + N} \geq \beta\,.$$

Theorem 2.1. *For any fixed positive even integer α, given a set $\mathcal{S}$ of transmitters (all of power 1) and a set $\mathcal{Q}$ of receivers, of sizes n and m respectively, we can determine which, if any, transmitter, is received by each receiver in total time $O((n+m)\log^2 n \log\log n)$.*

Proof. As pointed out above, a receiver q can receive only the closest transmitter $s(q)$, if any, as the SINR inequality implies $\frac{1}{|q-s(q)|^\alpha} > \frac{1}{|q-s|^\alpha}$ for any $s \neq s(q)$, or equivalently, $|q - s(q)| < |q - s|$. So, as a first step, we identify the closest transmitter for each receiver, which can be done, for example, by sorting $\mathcal{S}$, and using binary search for each receiver, in total time $O((m+n)\log n)$. Moreover, we can compute the term $\frac{1}{|q-s(q)|^\alpha}$, for each $q \in Q$, in the same amount of time.

Observe that f is a sum of n low-degree fractional functions of a single real variable q, so according to Corollary A.4, we can now evaluate f on all points of $\mathcal{Q}$ simultaneously in time $O((n+m)\log^2 n \log\log n)$.

In $O(m)$ additional operations we can evaluate the expressions $E(q_1), \ldots, E(q_m)$ and determine for which receivers the SINR inequality holds, so that the signal is actually received.

Computing and evaluating the fraction dominates the computation cost, so the total running time is $O((n+m)\log^2 n \log\log n)$. □

2.2 Arbitrary power

We proceed in precisely the same manner, except now we need the help of Fact B.5 to handle the multiplicatively weighted Voronoi diagram on a line. For omitted details, see Section D.

Theorem 2.2. *For any fixed positive even integer α, given a set $\mathcal{S}$ of transmitters (not necessarily all of the same power) and a set $\mathcal{Q}$ of receivers, of sizes n and m respectively, we can determine which, if any, transmitter is received by each receiver in total time $O((n+m)\log^2 n \log\log n)$.*

3 Batched point location in the plane

In this section $\mathcal{S} = \{s_i\}$ is a set of n point transmitters in the plane. We consider three versions of (batched) point location, where in the first two the answers we obtain are exactly correct, while in the third one the answer to a query q may be either "s" (meaning that q receives s), "no" (meaning that q does not receive any transmitter), or "maybe" (meaning that q may or may not be receiving some transmitter; the SIN ratio is too close to β and we are unable to decide quickly whether it is above or below β).

Specifically, we consider the following three versions of (batched) point location. In the first version, we assume that the *transmitters* form an $\sqrt{n} \times \sqrt{n}$ non-uniform grid and that each transmitter has power 1. We show how to solve a *single* point-location query in this setting in $O(\sqrt{n}\log^2 n \log\log n)$ (rather than linear) time; refer to Section 3.2. In the second version (Section 3.3), we assume that the *receivers* form an $n \times n$ non-uniform grid, but the n transmitters, on the other hand, are located anywhere in the plane. Moreover, we allow arbitrary transmission powers. We show how to answer the n^2 queries in near-quadratic (rather than cubic) time.

Finally, in the third version (Section 3.4), we do not make any assumptions on the location of the devices (either transmitters or receivers). As a result of this, we might not be able to give a definite answer in borderline instances. Specifically, given n transmitters and m receivers, we compute (in total time near-linear in $n+m$), for each receiver q, its unique candidate transmitter s and a value $\tilde{E}(q)$, such that, if $\tilde{E}(q)$ is sufficiently greater than β, then q surely receives s, if $\tilde{E}(q)$ is sufficiently smaller than β, then q surely does not receive s, and otherwise, q may or may not receive some transmitter (i.e., $\tilde{E}(q)$ lies in the *uncertainty interval*). We first present a solution for which the uncertainty interval is $[2^{-\alpha/2}\beta, 2^{\alpha/2}\beta)$, i.e., a constant-factor approximation. We then

generalize it so that the uncertainty region is $[(1-\varepsilon)\beta, (1+\varepsilon)\beta)$, for any $\varepsilon > 0$, i.e., a PTAS. We consider both the uniform- and arbitrary-power settings.

3.1 General discussion

Once again, the SINR inequality determines which, if any, of the transmitters $s \in S$ can be heard by a receiver at point q and the only candidate transmitter $s(q)$ is the one that minimizes $|q-s|/p^{1/\alpha}$ among all transmitters s with corresponding power p. In the uniform-power case, this corresponds to the transmitter closest to q in Euclidean distance, and the corresponding space decomposition is the Euclidean Voronoi diagram which can be constructed in $O(n \log n)$ time (see Section B.2), where $n = |\mathcal{S}|$. In the non-uniform-power case, this corresponds to the multiplicatively weighted Voronoi diagram in the plane, which is a structure of worst-case complexity $\Theta(n^2)$ that can be constructed in time $O(n^2)$.

Once again we define the function $f(q)$, which represents the total signal strength at q from *all* transmitters, and express the decision of whether the transmitter $s(q)$ is received at q by computing $E(q)$ and comparing it with β. The difference from the one-dimensional case is that $f(q)$ is now a sum of low-degree *bivariate* fractions, with the two variables being the coordinates of q.

In all cases, the goal is to evaluate $f(q)$, for each receiver q, and to identify the suitable candidate transmitter $s(q)$, faster than by brute force. Given this information, the decision can be made in constant time per receiver.

We now discuss each two-dimensional problem in turn.

3.2 Transmitters on a grid

In this version we assume that the transmitters of $\mathcal{S}$ form an $\sqrt{n} \times \sqrt{n}$ non-uniform grid and have uniform power. We show how to answer a *single* arbitrary point-location query in roughly $\sqrt{n}$ (rather than linear) time.

We assume that $\mathcal{S} = X \times Y$, where X and Y are two sets of $\sqrt{n}$ numbers (coordinates) each. We start by sorting X and Y.

As mentioned above, the problem reduces to computing $f(q)$ and identifying $s(q)$ for the query q. The latter is not difficult: since all powers are the same, it is sufficient to identify the point in $X \times Y$ closest to q. It is easy to check that the closest point is always one of the (at most) four corners of the grid cell containing q and can be found by one binary search on X and one on Y and comparing each of the at most four corresponding distances to q.

So we focus on the computation of $f(q)$. We rewrite it as

$$f(q; X; Y) = \sum_{x \in X} \sum_{y \in Y} \frac{1}{|q-(x,y)|^\alpha} = \sum_{x \in X} g_Y(q; x),$$

where we view

$$g_Y(q; x) \coloneqq \sum_{y \in Y} \frac{1}{|q-(x,y)|^\alpha}$$

as a sum of fractional functions of x only, having substituted the actual values for the numbers of Y. Using Corollary A.4, we can sum these fractions and evaluate the sum at the $\sqrt{n}$ distinct points $x_1, x_2, \ldots$ of X, obtaining $g_Y(q, x_1), g_Y(q, x_2), \ldots$ and thereby $f(q)$, in $O(\sqrt{n} \log^2 n \log\log n)$ operations, which dominates the running time.

Theorem 3.1. *For any fixed positive even integer α, given a set $\mathcal{S}$ of n transmitters (all of power 1) forming a $\sqrt{n} \times \sqrt{n}$ non-uniform grid and a receiver q, we can determine which, if any, transmitter, is received by q in time $O(\sqrt{n} \log^2 n \log\log n)$.*

Note. For a $k \times l$ grid, where $l \leq k$, this would take $O((k+l)\log^2 l \log\log l) + O((k+l)\log(k+l)) = O((k+l)\log^2 l \log\log l)$ time.

3.3 Receivers on a grid

In this version we assume that the receivers of $\mathcal{Q}$ form an $n \times n$ non-uniform grid. The n transmitters, on the other hand, are located anywhere in the plane. Moreover, we allow arbitrary transmission powers. We show how to answer the n^2 queries in near-quadratic (rather than cubic) time.

In this case, we need to evaluate $f(q)$, which is a sum of n bivariate low-degree fractions at all points q_{ij} of a two-dimensional (possibly non-uniform) grid $X \times Y$.[2] This is taken care of in $O(n^2 \log n \log\log n)$ time by Corollary A.8.

The only missing ingredient is identifying the candidate transmitter $s(q)$ for each $q \in \mathcal{Q}$. This can be done, for example, by computing the weighted Voronoi diagram, preprocessing it for point location, and executing n^2 such queries, one for each $q \in \mathcal{Q}$.

It turns out that the following alternative may be simpler to implement: Observe that on each line $x = x_j$, the functions $|q - s_i|/p_i^{1/\alpha}$ behave very similarly to the univariate case in Section 2.2 above, so the problem can be solved in $O(n \log n)$ time per line, for a total of $O(n^2 \log n)$ time; refer to Fact B.6.

Combining the computation of $f(q)$ over all points of $X \times Y$ and identification of $s(q)$, we obtain the claimed result in $O(n^2 \log n \log\log n)$ time.

Theorem 3.2. *For any fixed positive even integer α, given a set $\mathcal{S}$ of n transmitters (not necessarily all of the same power), and a set $\mathcal{Q}$ of n^2 receivers forming an $n \times n$ non-uniform grid, we can determine which, if any, transmitter, is received by each receiver in total time $O(n^2 \log n \log\log n)$.*

Note. There are two obstacles to speeding this up for a smaller number of receivers: (a) An explicit representation of $f(q)$ has $\Theta(n^2)$ coefficients and thus cannot be processed in subquadratic time. (b) It seems harder to identify the candidate transmitter in the non-uniform power case; the first solution builds a quadratic-size space decomposition and the second requires linear time for each set of collinear points; both result in near-quadratic performance. A new idea is needed.

Note. If we do not restrict the location of the n^2 receivers, we can still obtain a subcubic solution by applying a different algebraic technique. Indeed, according to Fact A.9, we can determine which, if any, transmitter, is received by each receiver in total time $O(n^{1+\omega_2/2+\varepsilon})$, for any $\varepsilon > 0$, where $\omega_2 < 3.334$ is a constant related to the efficiency of matrix multiplication.

3.4 Approximating the general case

We now abandon the ambition to get exact answers and aim for an approximation algorithm, in the sense we will make precise below. Again, $\mathcal{S} = \{s_i\}$ is the set of n transmitters, with each s_i a point in the plane with power p_i; similarly $\mathcal{Q} = \{q_j\}$ is the set of m receivers, where a generic receiver is $q = (q_x, q_y)$.

[2] Slanted and sheared grids can be handled by a simple extension of this observation; we omit the easy details.

For a query point q and a transmitter $s = (s_x, s_y)$ of power p, set $l(q,s) = \max\{|q_x - s_x|, |q_y - s_y|\}$; in other words, $l(q,s)$ is the L^∞ distance between points q and s. In complete analogy to our previous approach, put

$$\tilde{f}(q) \coloneqq \sum_{i=1}^{n} \frac{p_i}{l(q,s_i)^\alpha} \quad \text{and} \quad \tilde{E}(q) \coloneqq \frac{\frac{p}{l(q,s)^\alpha}}{\tilde{f}(q) - \frac{p}{l(q,s)^\alpha} + N}.$$

What is the significance of the quantity $\tilde{E}(q)$? Since for any two points s, q, $l(q,s) \le |q-s| \le \sqrt{2}l(q,s)$,

$$2^{-\alpha/2}\frac{p_j}{l(q,s_j)^\alpha} \le \frac{p_j}{|q-s_j|^\alpha} \le \frac{p_j}{l(q,s_j)^\alpha},$$

so $2^{-\alpha/2}\tilde{f}(q) \le f(q) \le \tilde{f}(q)$, and therefore $2^{-\alpha/2}\tilde{E}(q) \le E(q) \le 2^{\alpha/2}\tilde{E}(q)$. Informally, $\tilde{E}(q)$ is "pretty close" to $E(q)$.

This suggests an approximation strategy that begins by computing $\tilde{E}(q)$ instead of $E(q)$. If $\tilde{E}(q) \ge 2^{\alpha/2}\beta$, we know that $E(q) \ge \beta$ and the signal from the unique candidate transmitter $s(q)$ *is* received. If $\tilde{E}(q) < 2^{-\alpha/2}\beta$, then $E(q) < \beta$ and the signal from $s(q)$ is *not* received and therefore no signal is received by q. For intermediate values of $\tilde{E}(q)$, we cannot definitely determine whether $s(q)$'s signal is received at q.

Now we turn to the actual batch computation of $\tilde{E}(q)$ for all receivers in $\mathcal{Q}$ and point out a few additional caveats.

Computationally, $\tilde{E}(q)$ can be evaluated in constant time, given $\tilde{f}(q)$ and point $s(q) = s(q,\mathcal{S})$. So we focus on these two subproblems. For the uniform-power case, we can construct the Voronoi diagram of $\mathcal{S}$, preprocess it for point location, and query it with each receiver, for a total cost of $O((n+m)\log n)$; see Fact B.3. In the case of non-uniform power, if we are content with near-quadratic running time, we can determine $s(q)$ by computing the multiplicatively weighted Voronoi diagram of $\mathcal{S}$ as outlined above, and then querying it with each receiver in total time $O(n^2 + m\log n)$ (see Fact B.4), which is too much for $m \approx n$. We provide an alternative below.

We show how to compute the values $\tilde{f}(q_1), \ldots, \tilde{f}(q_m)$ in near-linear time, using a two-dimensional orthogonal range search tree. Indeed, observe that $l(s,q) = |q_x - s_x|$ provided $|q_x - s_x| \ge |q_y - s_y|$. For a fixed q, the region W_q containing the transmitters of $\mathcal{S}$ satisfying this inequality is a 90° double wedge. Using (a tilted version of) the orthogonal range search tree (see Section B.1), we can construct a pair decomposition $\{(\mathcal{S}_i, \mathcal{Q}_i)\}$ of small size, so that each pair (s,q) with $s \in W_q$ appears in exactly one product $\mathcal{S}_i \times \mathcal{Q}_i$.

We now denote by $\tilde{f}(q,Z)$ the sum analogous to $\tilde{f}(q)$, where the summation goes over the elements of the supplied set Z rather than those of $\mathcal{S}$. Clearly,

$$\tilde{f}(q, \mathcal{S} \cap W_q) = \sum_{i : q \in \mathcal{Q}_i} \tilde{f}(q, \mathcal{S}_i), \tag{1}$$

by the definition of the pair decomposition. The number of terms in the last sum is $O(\log^2 n)$. Notice that $\tilde{f}(q,\mathcal{S}_i)$, for a fixed i, is a sum of small fractional *univariate* functions, with $|\mathcal{S}_i|$ terms in it, since the expression for transmitters in W_q depends only on q_x and not on q_y. Now for each pair $(\mathcal{Q}_i, \mathcal{S}_i)$, we use Corollary A.4 to evaluate $\tilde{f}(q,\mathcal{S}_i)$ on each $q \in \mathcal{Q}_i$ in total time $O((|\mathcal{Q}_i| + |\mathcal{S}_i|)\log^2 |\mathcal{S}_i| \log\log |\mathcal{S}_i|) = O((|\mathcal{Q}_i| + |\mathcal{S}_i|)\log^2 n \log\log n)$. This gives us all the summands of (1) and therefore allows us to evaluate $\tilde{f}(q, \mathcal{S} \cap W_q)$ for all $q \in \mathcal{Q}$, in total time at most proportional to $\sum_i (|\mathcal{Q}_i| + |\mathcal{S}_i|)\log^2 n \log\log n = (\sum_i (|\mathcal{Q}_i| + |\mathcal{S}_i|))\log^2 n \log\log n = O((m+n)\log^4 n \log\log n)$.

Of course, we have only treated those s that lie in W_q. But the calculation is repeated in the complementary double wedge, where now only the y-coordinates matter and $\tilde{f}(q)$ is the sum of the two values thus obtained.

Theorem 3.3. *For any fixed positive even integer α, given a set $\mathcal{S}$ of n transmitters (all of power 1) and a set $\mathcal{Q}$ of m receivers, we can do the following in total time $O((m+n)\log^4 n \log\log n)$. For each $q \in \mathcal{Q}$, we find its unique candidate transmitter $s(q)$ and compute a value $\tilde{E}(q)$, such that (i) if $\tilde{E}(q) \geq 2^{\alpha/2}\beta$, then q can definitely hear $s(q)$, (ii) if $\tilde{E}(q) < 2^{-\alpha/2}\beta$, then q definitely cannot hear $s(q)$, and (iii) if $2^{-\alpha/2}\beta \leq \tilde{E}(q) < 2^{\alpha/2}\beta$, then q may or may not hear $s(q)$.*

The algorithm for the non-uniform power case is hampered by the fact that the obvious way to identify the candidate transmitter each receiver might hear seems to involve constructing the multiplicatively weighted Voronoi diagram of quadratic complexity. However, we do not need the exact multiplicatively closest neighbor, but rather a reasonably-close approximation of the value $|q-s|/p(s)^{1/\alpha}$, over all $s \in \mathcal{S}$ (being off by a multiplicative factor of at most $2^{1/2}$ is sufficient; see the discussion below). Such an approximation is provided by the first algorithm in Fact B.7, for a constant value of the approximation parameter ε (namely, $\varepsilon = 2^{1/2}-1$), yielding the following:

Theorem 3.4. *For any fixed positive even integer α and any $\beta > 2^{\alpha/2}$, given a set $\mathcal{S}$ of n transmitters of arbitrary powers and a set $\mathcal{Q}$ of m receivers, we can do the following in total time $O(n\log^7 n + m\log^4 n \log\log n)$ and $O(n\log^4 n)$ space: For each $q \in \mathcal{Q}$, we find a transmitter s_q and compute a value $\tilde{E}(q)$, such that (i) if $\tilde{E}(q) \geq 2^{\alpha/2}\beta$, then q can definitely hear s_q (implying that $s_q = s(q)$), (ii) if $\tilde{E}(q) < 2^{-\alpha/2}\beta$, then q definitely cannot hear any transmitter, and (iii) if $2^{-\alpha/2}\beta \leq \tilde{E}(q) < 2^{\alpha/2}\beta$, then q may or may not hear one of the transmitters.*

Note. The transmitter s_q in the theorem above is not necessarily the unique candidate transmitter $s(q)$. We would like to show that if $\tilde{E}(q) \geq 2^{\alpha/2}\beta$ (and therefore $E(q) \geq \beta$), then s_q is necessarily $s(q)$. Assume that they are different (i.e., that $s_q \neq s(q)$), and let e_q (resp., $e(q)$) be the strength of s_q's signal (resp., $s(q)$'s signal) at q. Then, we know that $e_q \leq e(q) \leq 2^{\alpha/2}e_q$. Notice that $E(q) \leq e(q)/e_q$, since $E(q)$ is maximized when there is no third transmitter and no noise, so $e(q)/e_q \geq \beta$ (since $E(q) \geq \beta$). Recall that we are assuming that $\beta > 2^{\alpha/2}$, so we get that $e(q)/e_q > 2^{\alpha/2}$, which is a contradiction.

We now turn the algorithm described above into a PTAS, in the sense that we will confine $\tilde{E}(q)$ to the range $((1-\varepsilon)E(q), (1+\varepsilon)E(q)]$, for a given $\varepsilon > 0$. We outline the approach below. Consider the regular k-gon K_k circumscribed around the Euclidean unit disk, for a large enough even $k \geq 4$ specified below. We modify the above algorithm, replacing the L^∞-norm whose "unit disk" is a square, with the norm $|\ldots|_k$ with K_k as the unit disk. Then $|v|_k \leq |v| \leq (1+\Theta(k^{-2})|v|_k$, for any vector v in the plane. In the range-searching data structure, wedges with opening angle $\pi/2 = 2\pi/4$ are replaced by wedges with opening angle $2\pi/k$, and we need $k/2$ copies of the structure.

In terms of the quality of approximation, the factor $2^{\alpha/2} = (\sqrt{2})^\alpha$ is replaced by $(1+\Theta(k^{-2}))^\alpha \approx 1+\alpha\Theta(k^{-2})$. Hence to obtain an approximation factor of $1+\varepsilon$, we set $1+\varepsilon = 1+\alpha\Theta(k^{-2})$, or $k = c(\alpha/\varepsilon)^{1/2}$, for a suitable absolute constant c. In other words, it is sufficient to create $O(\varepsilon^{-2})$ copies of the data structure. To summarize, we have:

Theorem 3.5. *For a positive ε, any fixed positive even integer α, given a set $\mathcal{S}$ of n transmitters (all of power 1) and a set $\mathcal{Q}$ of m receivers, we can do the following in total time $O((m+n)\varepsilon^{-2}\log^4 n \log\log n)$. For each $q \in \mathcal{Q}$, we find its unique candidate transmitter $s(q)$*

and compute a value $\tilde{E}(q)$, such that (i) if $\tilde{E}(q) \geq (1+\varepsilon)\beta$, then q can definitely hear $s(q)$, (ii) if $\tilde{E}(q) < (1-\varepsilon)\beta$, then q definitely cannot hear $s(q)$, and (iii) if $(1-\varepsilon)\beta \leq \tilde{E}(q) < (1+\varepsilon)\beta$, then q may or may not hear $s(q)$.

Theorem 3.6. *For a positive ε, any fixed positive even integer α, and any $\beta > 1+\varepsilon$,[3], given a set $\mathcal{S}$ of n transmitters of arbitrary powers and a set $\mathcal{Q}$ of m receivers, we can do the following in total time $O(n\varepsilon^{-8}\log^7 n + m\varepsilon^{-5}\log^4 n \log\log n)$ and $O(n\varepsilon^{-6}\log^4 n)$ space: For each $q \in \mathcal{Q}$, we find a transmitter s_q and compute a value $\tilde{E}(q)$, such that (i) if $\tilde{E}(q) \geq (1+\varepsilon)\beta$, then q can definitely hear s_q (implying that $s_q = s(q)$), (ii) if $\tilde{E}(q) < (1-\varepsilon)\beta$, then q definitely cannot hear any transmitter, and (iii) if $(1-\varepsilon)\beta \leq \tilde{E}(q) < (1+\varepsilon)\beta$, then q may or may not hear one of the transmitters.*

4 Concluding remarks

We described several algorithms that combine computational geometry techniques and methods of computer algebra to obtain very fast batched SINR diagram point-location queries.

We believe that Theorems 3.5 and 3.6 can be applied to speed up the preprocessing stage of existing point-location results. Consider, e.g., the data structure presented by Avin et al. [6] for a set of n uniform-power transmitters, whose construction time is $O(n^2/\delta)$. This data structure is actually a collection of n data structures, one per transmitter, where the data structure DS_i for transmitter s_i consists of an inner (R_i^+) and outer (R_i^-) approximation for reception region R_i, so that $area(R_i^- \setminus R_i^+) \leq \delta \cdot area(R_i)$, see the definitions in the introduction. The construction of DS_i is based on the convexity and fatness of region R_i and consists of two stages. In the first, explicit estimates for the radii of the largest disk centered at s_i and contained in R_i and the smallest such disk containing R_i are obtained, by applying a binary-search-like procedure (beginning with the distance between s_i to its nearest (other) transmitter in $\mathcal{S}$), where each comparison is resolved by explicitly evaluating the SIN ratio at some point q and comparing it to β, i.e., by an *in/out* test. In the second stage, a $1/\delta \times 1/\delta$ grid scaled to exactly cover the outer disk is laid, and, by performing $O(1/\delta)$ additional in/out tests, the sets R_i^+ and R_i^- are obtained (as collections of grid cells). This algorithm thus performs $\Theta(\log n + 1/\delta)$ in/out tests per transmitter, at a cost of $\Theta(n)$ operations each; the high cost of each test is the bottleneck.

We believe that it is possible to speed up the algorithm by constructing the n individual data structures in parallel. During the construction, we will form $O(\log n + 1/\delta)$ batches of n queries each, and use Theorem 3.5 to deal with each of them in near-linear time. The only problem is that our query answers are not exact, but approximate; for some queries, instead of "in" or "out," we answer "maybe. We think that there is a way to overcome this problem, but we leave it for a full version.

Besides speeding up the construction time of known structures, we would like to find other applications of batched point location to other problems studied in the SINR model.

We note that our results are general, in the sense that analogous results can be obtained for diagrams that are induced by other inequalities similar to the SINR inequality.

Finally, on a larger scale, we are interested in further applications where algebraic and geometric tools can be combined to achieve significant improvements.

Acknowledgments B.A. would like to acknowledge the help of Sariel Har-Peled in matters of approximation and of Guillaume Moroz in matters of algebra. He would also like to thank Pankaj K. Agarwal for general encouragement.

[3]This requirement is analogous to that in Theorem 3.4 to guarantee that the approximately highest-strength transmitter returned by the data structure is in fact the right one.

References

[1] D. Ajwani, S. Ray, R. Seidel, and H. R. Tiwary. On computing the centroid of the vertices of an arrangement and related problems. In *Algorithms and Data Structures, 10th International Workshop, WADS 2007, Halifax, Canada, August 15-17, 2007, Proceedings*, pages 519–528, 2007.

[2] M. Andrews and M. Dinitz. Maximizing capacity in arbitrary wireless networks in the SINR model: Complexity and game theory. In *INFOCOM*, pages 1332–1340, 2009.

[3] F. Aurenhammer. The one-dimensional weighted Voronoi diagram. *Inf. Process. Lett.*, pages 119–123, 1986.

[4] F. Aurenhammer and H. Edelsbrunner. An optimal algorithm for constructing the weighted Voronoi diagram in the plane. *Pattern Recognition*, pages 251–257, 1984.

[5] C. Avin, A. Cohen, Y. Haddad, E. Kantor, Z. Lotker, M. Parter, , and D. Peleg. SINR diagram with interference cancellation. In *Proceedings of Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 502–515, 2012.

[6] C. Avin, Y. Emek, E. Kantor, Z. Lotker, D. Peleg, and R. Liam. SINR diagrams: Convexity and its applications in wireless networks. *J. ACM*, 59(4):18:1–18:34, Aug. 2012.

[7] D. Bini and V. Y. Pan. *Polynomial and Matrix Computations (Vol. 1): Fundamental Algorithms*. Birkhauser Verlag, Basel, Switzerland, Switzerland, 1994.

[8] H. Edelsbrunner and R. Seidel. Voronoi diagrams and arrangements. *Discrete Comput. Geom.*, 1:25–44, 1986.

[9] O. Goussevskaia, M. M. Halldórsson, R. Wattenhofer, and E. Welzl. Capacity of arbitrary wireless networks. In *INFOCOM*, pages 1872–1880, 2009.

[10] O. Goussevskaia, Y. A. Oswald, and R. Wattenhofer. Complexity in geometric SINR. In *MobiHoc*, pages 100–109, 2007.

[11] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Trans. Information Theory*, 46(2):388–404, 2000.

[12] M. M. Halldórsson. Wireless scheduling with power control. *ACM Transactions on Algorithms*, 9(1), 2012.

[13] M. M. Halldórsson and P. Mitra. Wireless capacity with oblivious power in general metrics. In *SODA*, pages 1538–1548, 2011.

[14] M. M. Halldórsson and R. Wattenhofer. Wireless communication is in apx. In *ICALP*, pages 525–536, 2009.

[15] S. Har-Peled and N. Kumar. Approximating minimization diagrams and generalized proximity search. *SIAM J. Comput.* Accepted for publication; see `http://sarielhp.org/p/12/wann/wann.pdf`.

[16] E. Kantor, Z. Lotker, M. Parter, and D. Peleg. The topology of wireless communication. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 383–392, 2011.

[17] T. Kesselheim. A constant-factor approximation for wireless capacity maximization with power control in the SINR model. In *SODA*, pages 1549–1559, 2011.

[18] Z. Lotker and D. Peleg. Structure and algorithms in the SINR wireless model. *SIGACT News*, 41(2):74–84, 2010.

[19] G. Moroz and B. Aronov. Computing the distance between piecewise-linear bivariate functions. In *Proceedings of Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 288–293, 2012.

[20] G. Moroz and B. Aronov. Computing the distance between piecewise-linear bivariate functions. *ACM Transactions on Algorithms*, 2013. Accepted for publication; see also CoRR `abs/1107.2312`.

[21] T. Moscibroda and R. Wattenhofer. The complexity of connectivity in wireless networks. In *INFOCOM*, pages 23–29, 2006.

[22] M. Nüsken and M. Ziegler. Fast multipoint evaluation of bivariate polynomials. In *ESA*, pages 544–555, 2004.

[23] V. Y. Pan. Simple multivariate polynomial multiplication. *J. Symb. Comput.*, 18(3):183–186, 1994.

[24] M. Sharir and P. K. Agarwal. *Davenport-Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, New York, 1995.

[25] J. von zur Gathen. *Modern Computer Algebra*. Cambridge University Press, Cambridge, 1999.

[26] P.-J. Wan, X. Jia, and F. F. Yao. Maximum independent set of links under physical interference model. In *WASA*, pages 169–178, 2009.

A Tools and definitions from the algebra world

A.1 Definitions

A *univariate polynomial* $A(x)$ of degree at most n is defined by an expression of the form $A(x) = a_0 + a_1x + a_2x^2 + \cdots + a_nx^n$. The tuple $C(A) \coloneqq \langle a_0, a_1, \ldots, a_n\rangle$ is the *coefficient representation* of A. Given a set $X = \{x_0, \ldots, x_n\}$ of $n+1$ numbers, we let $V(A) = V(A, X)$ denote $\{(x_0, A(x_0)), \ldots, (x_n, A(x_n))\}$, a *value representation* of A.

A *fractional function* $F(x)$ of degree at most n is a function that can be written in the form $F(x) = A(x)/B(x)$ for two polynomials A and B of degree at most n, with $B \not\equiv 0$. *Coefficient representation* of F is simply $(C(A), C(B))$, the pair of coefficient representations of its numerator and denominator.

Analogously, a *bivariate polynomial* $B(x, y)$ of degree at most n in each variable is defined by an expression of the form $B(x, y) = \sum_{0\le i,j\le n} b_{ij}x^iy^j$. The set $\{\langle i, j, b_{ij}\rangle\}$ is the *coefficient representation* of B. The set $\{(x_i, y_i, B(x_i, y_i))\}$ is a *value representation* of B, for a suitably large set of points $\{(x_i, y_i)\} \in \mathbb{R}^2$.

One can analogously define bivariate fractional functions.

A.2 Univariate tools and facts

In this and the following sections we state several well-known results from computer algebra and refer to [25] for details; see [7] for an alternative presentation. The following two results can be found in Corollaries 10.8 and 10.12 in [25], respectively.

Fact A.1 (Univariate Multipoint Evaluation and Interpolation)**.** *If A is a univariate polynomial of degree at most n and X a tuple of $n+1$ distinct numbers, then the coefficient representation clearly determines A, but so does its value representation $V(A, X)$.*

Multipoint Evaluation *Given the coefficient representation $C(A)$ of A, A can be evaluated at the $n+1$ points of X, yielding $V(A, X)$ in $O(n\log^2 n\log\log n)$ arithmetic operations. If X contains m points, the computation can be done in $O((n+m)\log^2 n\log\log n))$ arithmetic operations.*[4]

Interpolation *Given a value representation $V(A, X)$ of a polynomial A of degree at most n on a set X of $n+1$ points, the coefficient representation $C(A)$ of A, can be constructed using $O(n\log^2 n\log\log n)$ arithmetic operations.*

Given two univariate polynomials A and B of degree at most n each, let their *product* be the polynomial D defined by $D(x) = A(x)B(x)$ for all x. Multiplying two polynomials in value form is easy: for each $x_i \in X$, $D(x_i) = A(x_i) \cdot B(x_i)$, by definition. Somewhat surprisingly, one can also quickly multiply univariate polynomials in coefficient representation (see Theorem 8.23 in [25]):

Fact A.2 (Fast Multiplication of Univariate Polynomials)**.** *Given two univariate polynomials A, B of degree at most n each, in coefficient representation, one can construct their product AB, using $O(n\log n\log\log n)$ arithmetic operations.*

[4] The second statement follows immediately from the first by extending X to length n with dummy data if $m < n$, or cutting it into chunks of length n if $m > n$.

Note. The Fast Fourier Transform (FFT) is essentially an evaluation of a polynomial on a special set of complex numbers, the nth roots of unity (more precisely, the $2^{\lceil \log_2 n \rceil}$th roots of unity). It can be performed in $O(n \log n \log\log n)$ operations, due to the special structure of the set of roots of unity; technically, the $\log\log n$ term appearing in many of our bounds is due to the assumption that the appropriate primitive root of unity is not available explicitly; otherwise the bounds improve by a factor of $\log\log n$. The Inverse FFT reverses the process (by using a variant of the FFT code), reconstructing a polynomial from its values at the roots of unity, again in $O(n \log n \log\log n)$ operations. The fast polynomial multiplication algorithm is an FFT, followed by point-wise multiplication, followed by the Inverse FFT.

The following observation has been made by previous authors; we provide a proof for completeness; we follow the construction from [19, 20].

Lemma A.3 (Sum of Fractions [19, 20])**.** *Given a set of n fractional functions $P_i(x)/Q_i(x)$ of constant degree each, in coefficient representation, the coefficient representation of their sum can be constructed using $O(n \log^2 n \log\log n)$ arithmetic operations.*

Proof. Given two fractions $\frac{A}{B}$ and $\frac{C}{D}$, their sum can be written as $\frac{AD+BC}{BD}$. Therefore, given two fractions of degree at most d in coefficient representation, we can obtain their sum, which is a fraction of degree at most $2d$, in coefficient representation, using three polynomial multiplications and one addition, for a total of $O(d \log d \log\log d)$ arithmetic operations (in fact, one can do better, though it does not affect asymptotic efficiency: convert all four polynomials into value form; construct the value form of the numerator and denominator of the sum; only then convert each of them back to coefficient form).

We start with the initial fractions P_i/Q_i, $i = 1, \ldots, n$ and add them in pairs, then add the resulting sums in pairs, and so forth. One round costs $O(n \log n \log\log n)$ operations, since the sum of the degrees of intermediate polynomials at each level is $O(n)$. In each round, the number of fractions reduces by a factor of two, so the procedure stops after $O(\log n)$ rounds. Hence the final result will be obtained after $O(n \log^2 n \log\log n)$ operations. □

Corollary A.4 (Evaluation of Sum of Fractions)**.** *Given a set of n fractional functions $P_i(x)/Q_i(x)$ of constant degree each, in coefficient representation, and a set of m values x_j, one can compute the m values $\sum_i P_i(x_j)/Q_i(x_j)$ in time $O((n+m) \log^2 n \log\log n)$.*

Proof. Explicitly sum the fractions to construct one degree-$O(n)$ fraction, using Lemma A.3, evaluate the numerator and denominator at each x_j using Fact A.1, and divide. □

A.3 Bivariate tools and facts

It would be helpful to have analogous tools for the bivariate case. The difficulty is that a general bivariate polynomial of degree at most n in each variable is described by $(n+1)^2 = \Theta(n^2)$ coefficients, so its explicit coefficient representation necessarily has size $\Theta(n^2)$. Thus in general there would be no hope for the existence of direct analogs of Facts A.1 and A.2, and Lemma A.3, or at least not in the same form.

We define the *product* $D = A \cdot B$ of two bivariate polynomials A and B similarly to that of univariate ones. Once again, multiplying two polynomials in value form is easy by multiplying the corresponding values. Somewhat surprisingly, one can also quickly multiply bivariate polynomials in coefficient representation:

Fact A.5 (Fast Multiplication of Bivariate Polynomials; [7, 23] or [25, Corollary 8.28])**.** *Given two bivariate polynomials $A(x, y)$, $B(x, y)$ of degree at most n in each variable, in coordinate representation (which has size $\Theta(n^2)$ in general), one can construct their product AB, using $O(n^2 \log n \log\log n)$ arithmetic operations.*

Notice that, once we have the multiplication result, one can use the same reasoning as in Lemma A.3 to obtain the following:

Fact A.6 (Sum of Bivariate Fractions)**.** *Given a set of n bivariate fractional functions $P_i(x, y)/Q_i(x, Y)$ of constant degree each, in coefficient representation, the coefficient representation of their sum $P(x, y)/Q(x, y)$, can be constructed using $O(n^2 \log n \log\log n)$ arithmetic operations.*

Fact A.7 (Bivariate Evaluation and Interpolation over a Grid [7, 23])**.** *Consider a bivariate polynomial $A(x, y)$ of degree at most n in each variable, and two sets $X = \{x_i\}$ and $Y = \{y_j\}$ of n numbers each.*

Grid Evaluation *Given the $O(n^2)$ coefficients of A, A can be evaluated at the points of $X \times Y$ (i.e., values $A(x_i, y_j)$ computed, for all combinations of i and j) in time $O(n^2 \log n \log\log n)$; see for example [7, 23].*

Interpolation from Values on the Grid *Given n^2 values $v_{ij} = A(x_i, y_j)$, one can reconstruct the coefficient representation of A in time $O(n^2 \log n \log\log n)$.*

Note. The evaluation bound described above is not very difficult to achieve: Write $A(x, y) = \sum_{j=0}^{n} B_j(x) y^j$, where B_j is a univariate polynomial of degree at most n and use the univariate algorithm from Fact A.1 to multi-evaluate B_j on all values x_i; the remaining arithmetic can be done in $O(n^2)$ operations.

Corollary A.8 (Evaluation of Sum of Fractions)**.** *Given a set of n bivariate fractional functions $P_i(x, y)/Q_i(x, y)$ of constant degree each, in coefficient representation and two sets $X = \{x_j\}$, $Y = \{y_k\}$ of n values each, one can compute the $n \times n$ values $\sum_i P_i(x_j, y_k)/Q_i(x_j, y_k)$ in time $O(n^2 \log n \log\log n)$.*

Proof. Combine the fractions into a single fraction using Fact A.6 and then evaluate on the grid $X \times Y$ using Fact A.7. □

There does not seem to be an easy analogue of fast univariate evaluation on an arbitrary set of values, as in Fact A.1, though a simple observation (see for example [22]) shows that one can evaluate a bivariate polynomial on n^2 points in roughly cubic time by just extending each set of n points into a grid and using the algorithm from Fact A.7. However, a stronger result is known

Fact A.9 (General Bivariate Multipoint Evaluation [22])**.** *If $A(x, y)$ is a bivariate polynomial of degree at most n in x and degree at most m in y, in coefficient representation (nm coefficients), A can be evaluated simultaneously at N different points of $\mathbb{R}^2$ in $O((N + nm)m^{\omega_2/2-1+\varepsilon})$ arithmetic operations, for any $\varepsilon > 0$, where ω_2 is such that the product of any $n \times n$ and $n \times n^2$ matrices can be performed in $O(n^{\omega_2})$ time; it is known that $\omega_2 < 3.334$.*

B Definitions and tools from geometry and data structures

B.1 Orthogonal range search trees

We say that point $p \in \mathbb{R}^2$ is *dominated* by point $q \in \mathbb{R}^2$ if $x(p) \leq x(q)$ and $y(p) \leq y(q)$; we write $p \leq q$.

Fact B.1. *Given a set $P = \{p_1, \ldots, p_n\}$ of n points in $\mathbb{R}^2$, a* (two-dimensional) orthogonal range tree *is a data structure of size $O(n \log n)$ that supports queries of the following type:*[5] *Given a query point q, report all the points $p \in P$ dominating q, in time $O(k + \log^2 n)$, where k is the size of the answer. More specifically, internally the range tree is a collection of $O(n \log n)$* canonical subsets *of P, so that each point of P lies in $O(\log^2 n)$ canonical subsets, and the answer of to the query q is represented as a disjoint union of $O(\log^2 n)$ such subsets.*

Fact B.2 (Dominating Pair Decomposition)**.** *Given two sets $P, Q \in \mathbb{R}^2$ each of size n, it is possible to construct a collection of pairs of subsets $\{(P_i, Q_i)\}$ with the following properties:*

1. *$P_i \subseteq P$, $Q_i \subseteq Q$.*
2. *For each pair of points $p \in P$, $q \in Q$ with $p \leq q$, there is a unique i so that $p \in P_i$ and $q \in Q_i$.*
3. *$\bigcup_i (P_i \times Q_i)$ contains only the dominating pairs.*
4. *There are at most $O(n \log n)$ pairs of sets in the collection.*
5. *$\sum_i (|P_i| + |Q_i|) = O(n \log^2 n)$.*
6. *The collection of pairs can be constructed in time $O(n \log^2 n)$.*

The above construction is not difficult if one starts with a range tree for P. The P_i are the canonical sets of the range tree. Initialize each Q_i to the empty set. One then executes a query for each $q \in Q$ and adds q to Q_i for each canonical set P_i that participates in the query.

B.2 Voronoi diagrams and friends

Given a set P of n points in the plane, one can partition the plane into convex polygonal regions, so that each region $V(p)$ consists entirely of points q closer to a given point $p \in P$ than to any other point of P, where "closer" is measured with respect to the usual Euclidean distance. Boundaries of the regions contain points where the nearest neighbor is non-unique. The partition is called the *Voronoi diagram* of P and the sets $V(p)$ are the Voronoi cells. The Voronoi diagram consists of n cells and a linear number of edges and vertices separating the cells.

Fact B.3 (Two-Dimensional Voronoi Diagram)**.** *Given an n-point set $P \subset \mathbb{R}^2$, the Voronoi diagram of P has linear complexity and can be constructed in time $O(n \log n)$ and preprocessed for logarithmic-time* point location *queries: Given a point q, determine the point $p \in P$ whose region $V(p)$ contains q.*

[5] A range tree is actually more powerful, but we will not need all of its power here.

If the Euclidean distance $|q-p|$ in the above construction is replaced by $|q-p|/w_p$, where w_p is a positive (multiplicative) weight, we obtain the *multiplicatively weighted Voronoi diagram*, which is similar to the Euclidean version, except that the complexity can be quadratic in n, in the worst case, the edges are circular arcs, and in general the region of a given Voronoi site can be disconnected and may contain holes.

Fact B.4 (Two-Dimensional Multiplicatively Weighted Voronoi Diagram [4])**.** *Given an n-point set $P \subset \mathbb{R}^2$ with positive weights w_p, the multiplicatively weighted Voronoi diagram of P has complexity $O(n^2)$ and can be constructed in time $O(n^2)$ and preprocessed for logarithmic-time* point-location queries*: Given a point q, determine the point $p \in P$ whose region $V(p)$ contains q.*

The one-dimensional analogs of the Euclidean and multiplicatively weighted Voronoi diagrams for an n-point set have linear complexity and can be constructed and preprocessed for point location in time $O(n \log n)$; the weighted version is specifically addressed in [3]:

Fact B.5 (One-Dimensional Multiplicatively Weighted Voronoi Diagram)**.** *Given an n-point set $P \subset \mathbb{R}$ with positive weights w_p, the multiplicatively weighted Voronoi diagram of P has linear complexity and can be constructed in time $O(n \log n)$ and preprocessed for logarithmic-time* point-location queries*: Given a point q, determine the point $p \in P$ whose region $V(p)$ contains q.*

We will also need the following slightly esoteric variant of the weighted diagram (which of course is a generalization of Fact B.5):

Fact B.6 (One-Dimensional Slice of a Two-Dimensional Multiplicatively Weighted Voronoi Diagram)**.** *Given an n-point set $P \subset \mathbb{R}$ with positive weights w_p, the multiplicatively weighted Voronoi diagram of P restricted to a given line ℓ has complexity $O(n)$ and can be constructed in time $O(n \log n)$ and preprocessed for logarithmic-time* point-location queries*: Given a point q, determine the point $p \in P$ whose region $V(p)$ contains q.*

Proof. Without loss of generality, after a suitable rigid transformation, we can assume that the line coincides with the x-axis. The squared distance function $f_p(x)$ from point $(x, 0)$ to site p is given by

$$f_p(x) = |p - (x,0)|^2/w_p^2 = ((x(p) - x)^2 + y(p)^2)/w_p^2,$$

which is a quadratic function of x. As observed, for example, by Edelsbrunner and Seidel [8], the minimization diagram of these functions coincides with the desired restricted weighted Voronoi diagram. Being quadratic, the graphs of two such functions $f_p(x)$ and $f_q(x)$ intersect at most twice and therefore their lower envelope and minimization diagram has linear complexity and can be constructed by a straightforward $O(n \log n)$ time algorithm [24, Theorem 6.1]. □

We will also need an approximate version of the Voronoi diagram. We will only state the weighted version of the problem, as that is what we need for our purposes [15, Theorem 2.16]:

Fact B.7 (Approximate Two-Dimensional Multiplicatively Weighted Nearest Neighbor [15])**.** *Given an n-point set $P \subset \mathbb{R}^2$ with positive weights w_p and a positive number ε, one can preprocess it into a data structure of space $O(n\varepsilon^{-6} \log^4 n)$ in time $O(n\varepsilon^{-6} \log^7 n)$ to support $O(\log n + 1/\varepsilon^3)$ time queries of the form: Given a point q, return $p' \in P$ so that $|p' - q|/w_{p'} \leq (1+\varepsilon)|p^* - q|/w_{p^*}$, where p^* is the point in P minimizing $|p - q|/w_p$.*

Alternatively, there is a data structure of space $O(n\varepsilon^{-6} \log^4 n)$ built in time $O(n\varepsilon^{-6} \log^7 n)$ with $O(\log(n/\varepsilon))$ query time.

C Handling directional antennas

In Section 3.4 we have been assuming that the transmitters in $\mathcal{S}$ are using omni-directional antennas. If directional antennas are used instead, the SINR inequality must be modified accordingly. The coverage area of the directional antenna of transmitter s is a wedge $W = W(s)$ of angle $\alpha = \alpha(s)$ and apex at s. Only points within W can receive s. Assuming that $q \in W$, we need to evaluate the left side of the SINR inequality, taking into account only the transmitters of $\mathcal{S}$ whose wedges cover q. That is, q receives s of power p if and only if

$$\frac{\frac{p}{|q-s|^\alpha}}{\sum_{s_j \neq s : q \in W_j} \frac{p_j}{|q-s_j|^\alpha} + N} \geq \beta \,.$$

The results of Section 3.4 are still relevant (as we show below), but they cannot be applied immediately. We first need to construct a data structure for range searching with wedges among the set of n points $\mathcal{Q}$: Given a query wedge $W = W(s)$, return $\mathcal{Q} \cap W$ as a union of a small number of canonical subsets of $\mathcal{Q}$. Then, we perform a query with each of the wedges $W(s)$, $s \in \mathcal{S}$, to obtain a collection of pairs of sets $\{(\mathcal{S}_i, \mathcal{Q}_i)\}$ with the property that, for any $(s, q) \in \mathcal{S} \times \mathcal{Q}$, if $q \in W(s)$, then there exists a single pair $(\mathcal{S}_i, \mathcal{Q}_i)$ such that $(s, q) \in \mathcal{S}_i \times \mathcal{Q}_i$, and, if $q \notin W(s)$, then $(s, q) \notin \mathcal{S}_i \times \mathcal{Q}_i$, for any i. We now can apply the results of the previous section to each of the pairs $(\mathcal{S}_i, \mathcal{Q}_i)$ to obtain theorems analogous to Theorems 3.3–3.6. The time bound in each of these theorems will be roughly $O(n^{4/3})$, assuming $|\mathcal{S}| = |\mathcal{Q}| = n$. We omit the remaining details from this version.

D Missing details: Arbitrary power transmitters on a line

Observe that, if q hears any of the transmitters in $\mathcal{S}$, then it is the *candidate* transmitter $s(q)$, for which the expression $p(s)/|q - s|^\alpha$ is maximum, over all $s \in \mathcal{S}$; $p(s)$ is the power of transmitter s. Thus, we find for each $q \in \mathcal{Q}$ its corresponding transmitter, by computing the appropriate weighted Voronoi diagram and querying it. This can be done in $O((n + m) \log n)$ time (see Fact B.5). Next, we generalize the function f defined above to mean

$$f(q) \coloneqq \sum_{j=1}^{n} \frac{p_j}{|q - s_j|^\alpha} \,.$$

Now, q hears $s(q)$ if and only if

$$E(q) \coloneqq \frac{\frac{p(s(q))}{|q-s(q)|^\alpha}}{f(q) - \frac{p(s(q))}{|q-s(q)|^\alpha} + N} \geq \beta \,,$$

and we can finish the proof of Theorem 2.2:

Proof of Theorem 2.2. Instead of looking for the nearest transmitter to q, we need to look for the transmitter maximizing $p_i/|q - s_i|^\alpha$, which is the same as minimizing $|q - s_i|/p_i^{1/\alpha}$. This corresponds to the one-dimensional *multiplicatively weighted Voronoi diagram* of $\mathcal{S}$ with weights $p_i^{1/\alpha}$, which, as per Fact B.5 and surrounding discussion, has complexity $O(n)$ and can be constructed in $O(n \log n)$ time; see [3] and [24, Theorem 6.1]. We once again use binary search to identify, for each receiver q_i,

the only possible transmitter $s(q_i)$ that it may be able to hear, compute the term $p(s(q_i))/|q_i - s(q_i)|^\alpha$ needed to finish computing $E(q_i)$, and thereby check the reception.

The rest of the algorithm proceeds as in Theorem 2.1: the function f depends on the power of individual transmitters, but it is still a sum of n low-degree univariate fractions and the remainder of the analysis follows unchanged. □